PROFESSIONAL-MACHINE-LEARNING-ENGINEER^{Q&As}

Professional Machine Learning Engineer

Pass Google PROFESSIONAL-MACHINE-LEARNING-ENGINEER Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.pass2lead.com/professional-machine-learning-engineer.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Google Official Exam Center

PROFESSIONAL-MACHINE-LEARNING-ENGINEER PDF Dumps |
PROFESSIONAL-MACHINE-LEARNING-ENGINEER VCE Dumps |
PROFESSIONAL-MACHINE-LEARNING-ENGINEER Exam Questions

1 / 5

![Pass2Lead logo](https://Pass2Lead.com)
⚙ **Instant Download** After Purchase

⚙ **100% Money Back** Guarantee

⚙ **365 Days** Free Update

⚙ **800,000+** Satisfied Customers

PROFESSIONAL-MACHINE-LEARNING-ENGINEER PDF Dumps |
PROFESSIONAL-MACHINE-LEARNING-ENGINEER VCE Dumps |
PROFESSIONAL-MACHINE-LEARNING-ENGINEER Exam Questions

2 / 5

![Pass2Lead](https://Pass2Lead.com)
QUESTION 1

You deployed an ML model into production a year ago. Every month, you collect all raw requests that were sent to your model prediction service during the previous month. You send a subset of these requests to a human labeling service to evaluate your model\'s performance. After a year, you notice that your model\'s performance sometimes degrades significantly after a month, while other times it takes several months to notice any decrease in performance. The labeling service is costly, but you also need to avoid large performance degradations. You want to determine how often you should retrain your model to maintain a high level of performance while minimizing cost. What should you do?
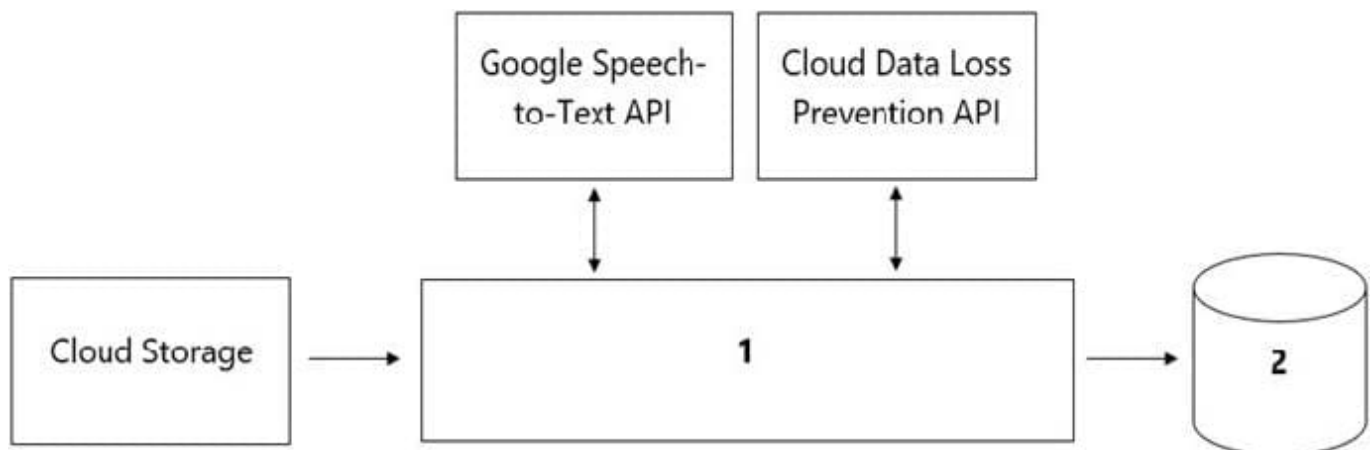
A. Train an anomaly detection model on the training dataset, and run all incoming requests through this model. If an anomaly is detected, send the most recent serving data to the labeling service.

B. Identify temporal patterns in your model\'s performance over the previous year. Based on these patterns, create a schedule for sending serving data to the labeling service for the next year.

C. Compare the cost of the labeling service with the lost revenue due to model performance degradation over the past year. If the lost revenue is greater than the cost of the labeling service, increase the frequency of model retraining; otherwise, decrease the model retraining frequency.

D. Run training-serving skew detection batch jobs every few days to compare the aggregate statistics of the features in the training dataset with recent serving data. If skew is detected, send the most recent serving data to the labeling service.

Correct Answer: D

https://cloud.google.com/blog/topics/developers-practitioners/monitor-models-training-serving-skew-vertex-aiew-vertex-aiandved=2ahUKEwiRg_aoj9n8AhWb7TgGHcGCDREQFnoECAwQAQandusg=AOvVaw197NneIJM0ra7fLq2zsOin

---

QUESTION 2

Your organization\'s call center has asked you to develop a model that analyzes customer sentiments in each call. The call center receives over one million calls daily, and data is stored in Cloud Storage. The data collected must not leave the region in which the call originated, and no Personally Identifiable Information (PII) can be stored or analyzed. The data science team has a third-party tool for visualization and access which requires a SQL ANSI-2011 compliant interface. You need to select components for data processing and for analytics. How should the data pipeline be designed?

PROFESSIONAL-MACHINE-LEARNING-ENGINEER PDF Dumps |
PROFESSIONAL-MACHINE-LEARNING-ENGINEER VCE Dumps |
PROFESSIONAL-MACHINE-LEARNING-ENGINEER Exam Questions

3 / 5

A. 1= Dataflow, 2= BigQuery

B. 1 = Pub/Sub, 2= Datastore

C. 1 = Dataflow, 2 = Cloud SQL

D. 1 = Cloud Function, 2= Cloud SQL

Correct Answer: A

https://github.com/GoogleCloudPlatform/dataflow-contact-center-speech-analysis

QUESTION 3

You are training an LSTM-based model on AI Platform to summarize text using the following job submission script:

gcloud ai-platform jobs submit training $JOB_NAME \

--package-path $TRAINER_PACKAGE_PATH \

--module-name $MAIN_TRAINER_MODULE \

--job-dir $JOB_DIR \

--region $REGION \

--scale-tier basic \

-- \

--epochs 20 \

--batch_size=32 \

--learning_rate=0.001 \

You want to ensure that training time is minimized without significantly compromising the accuracy of your model. What should you do?

A. Modify the `epochs\\' parameter.

B. Modify the `scale-tier\\' parameter.

C. Modify the `batch size\\' parameter.

D. Modify the `learning rate\\' parameter.

Correct Answer: B

Changing the scale tier does not impact performance?nly speeds up training time. Epochs, Batch size, and learning rate all are hyperparameters that might impact model accuracy.

QUESTION 4

PROFESSIONAL-MACHINE-LEARNING-ENGINEER PDF Dumps |
PROFESSIONAL-MACHINE-LEARNING-ENGINEER VCE Dumps |
PROFESSIONAL-MACHINE-LEARNING-ENGINEER Exam Questions

4 / 5

You need to build an ML model for a social media application to predict whether a user\'s submitted profile photo meets the requirements. The application will inform the user if the picture meets the requirements. How should you build a model to ensure that the application does not falsely accept a non-compliant picture?

A. Use AutoML to optimize the model\'s recall in order to minimize false negatives.

B. Use AutoML to optimize the model\'s F1 score in order to balance the accuracy of false positives and false negatives.

C. Use Vertex AI Workbench user-managed notebooks to build a custom model that has three times as many examples of pictures that meet the profile photo requirements.

D. Use Vertex AI Workbench user-managed notebooks to build a custom model that has three times as many examples of pictures that do not meet the profile photo requirements.

Correct Answer: B

QUESTION 5

You work for a large technology company that wants to modernize their contact center. You have been asked to develop a solution to classify incoming calls by product so that requests can be more quickly routed to the correct support team. You have already transcribed the calls using the Speech-to-Text API. You want to minimize data preprocessing and development time. How should you build the model?

A. Use the AI Platform Training built-in algorithms to create a custom model.

B. Use AutoMIL Natural Language to extract custom entities for classification.

C. Use the Cloud Natural Language API to extract custom entities for classification.

D. Build a custom model to identify the product keywords from the transcribed calls, and then run the keywords through a classification algorithm.

Correct Answer: B

PROFESSIONAL-MACHIN
E-LEARNING-ENGINEER
PDF Dumps

PROFESSIONAL-MACHIN
E-LEARNING-ENGINEER
VCE Dumps

PROFESSIONAL-MACHIN
E-LEARNING-ENGINEER
Exam Questions